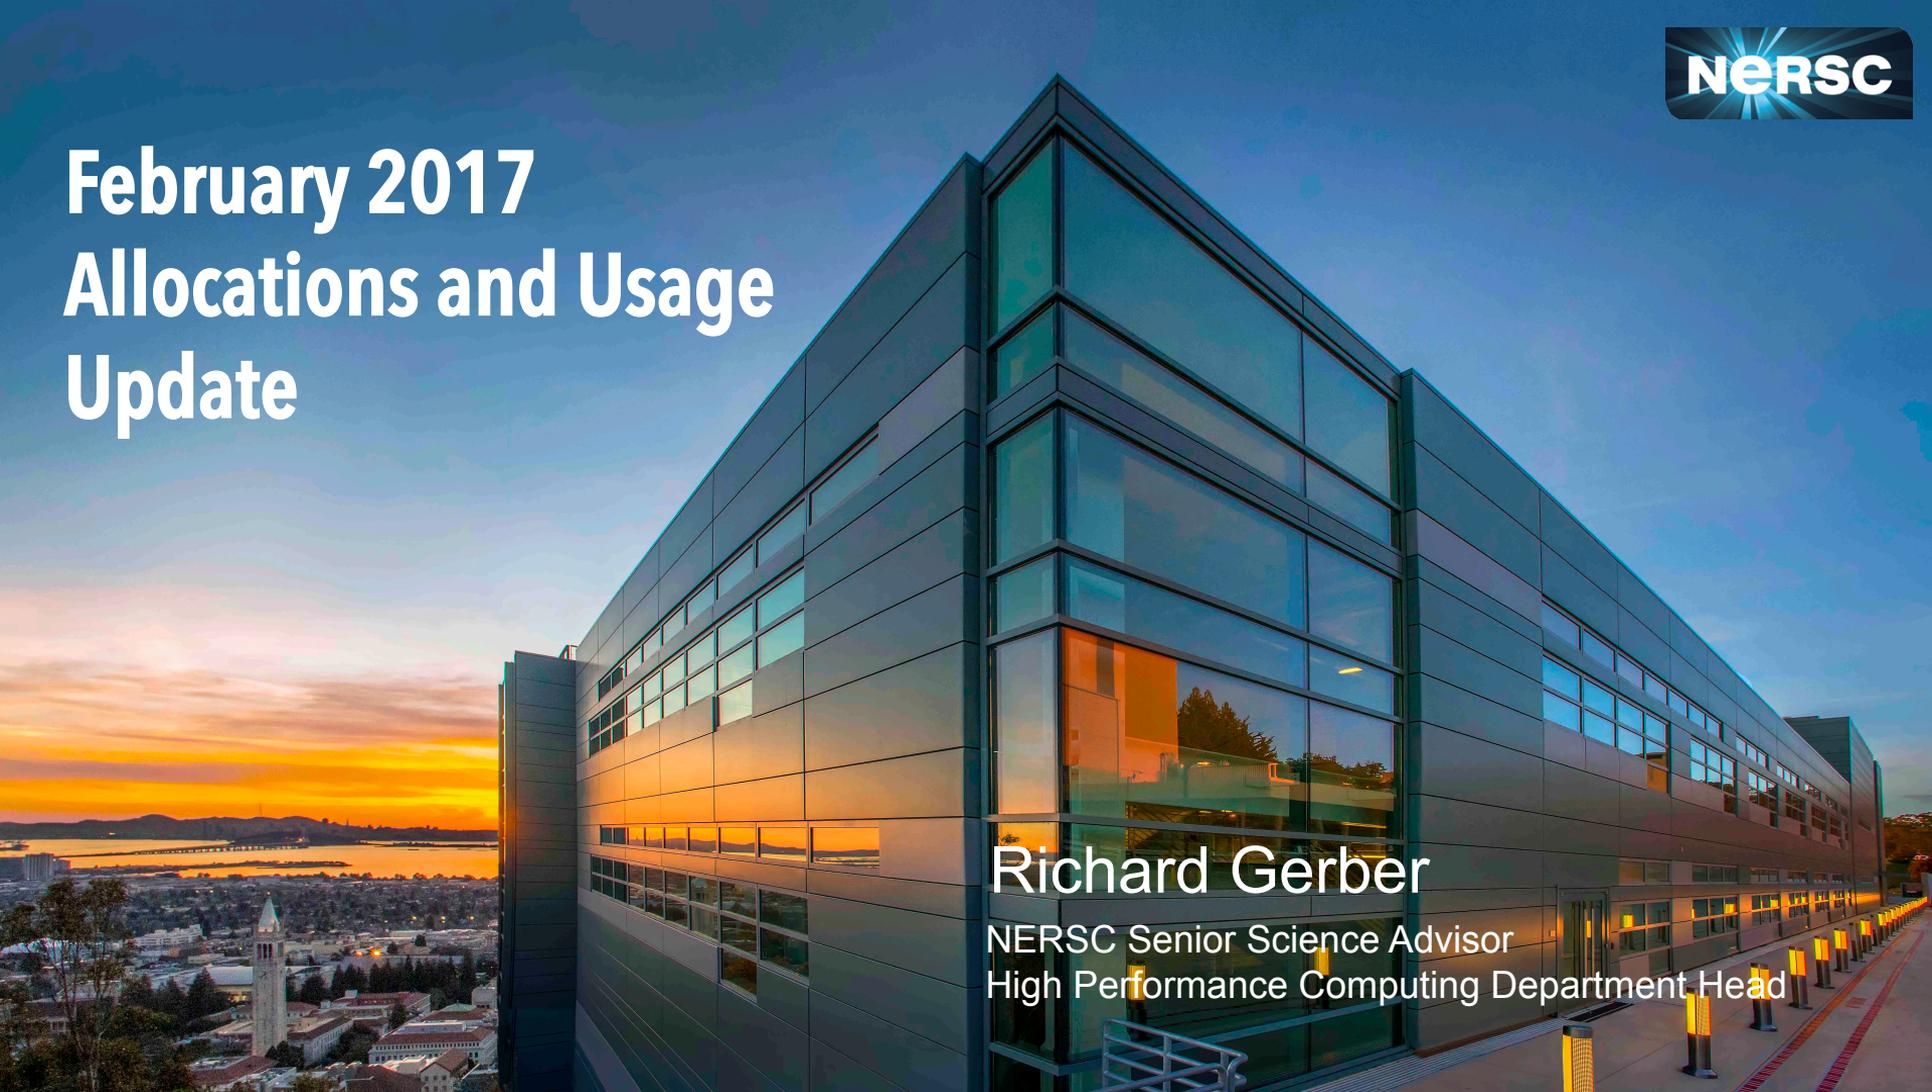


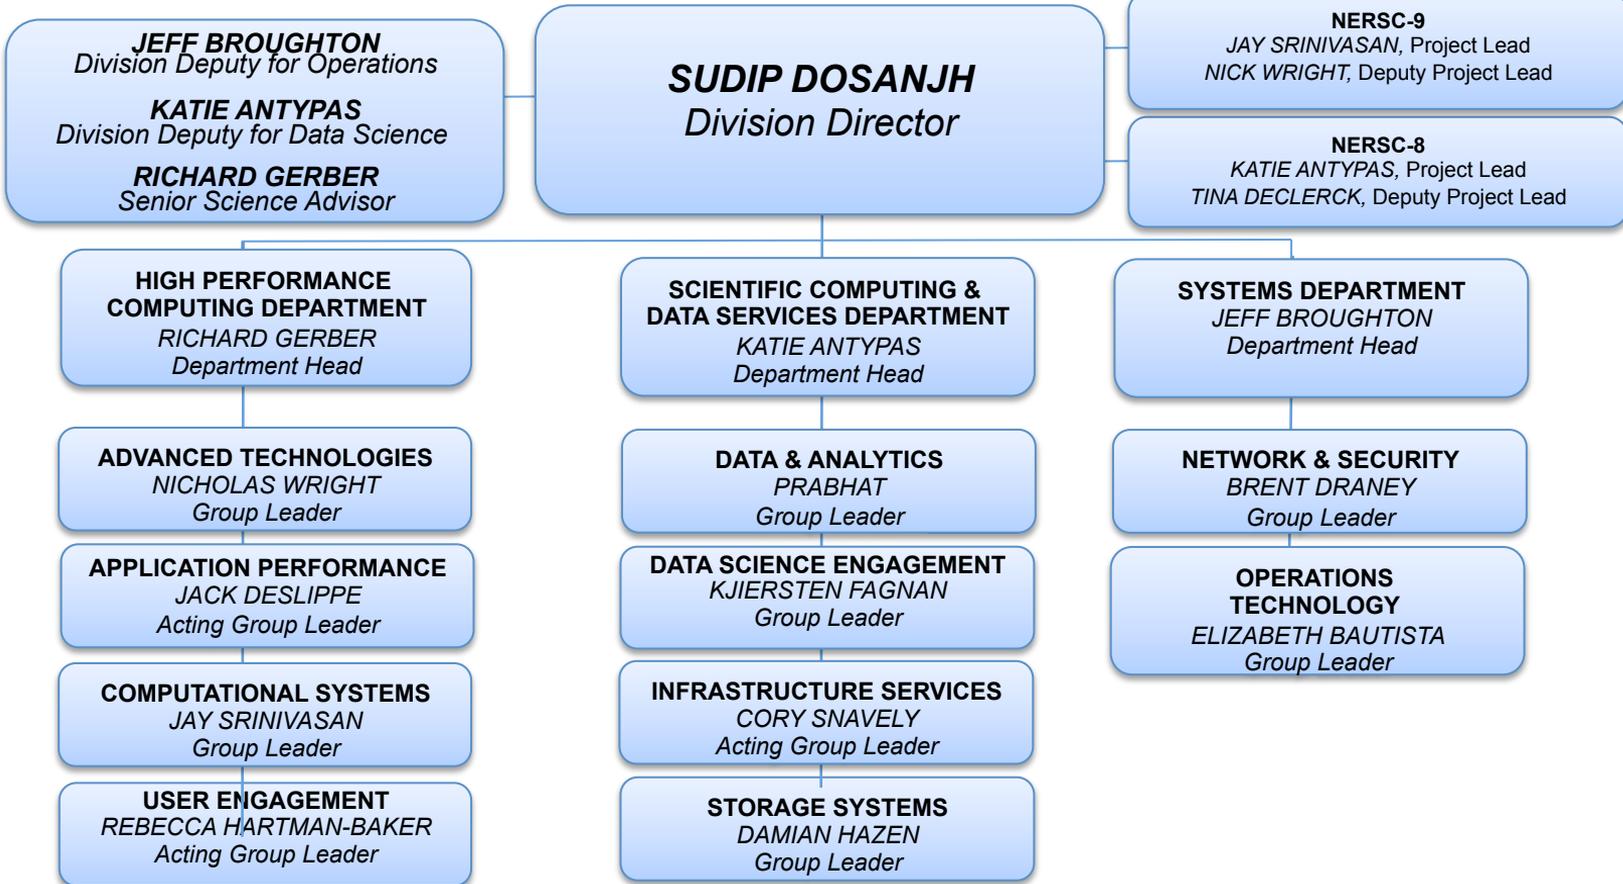
February 2017 Allocations and Usage Update

The background of the slide is a photograph of a modern, multi-story building with a glass and metal facade. The building is illuminated from within, and the windows reflect the sunset sky. In the background, a cityscape and a body of water are visible under a colorful sky transitioning from orange to blue.

Richard Gerber

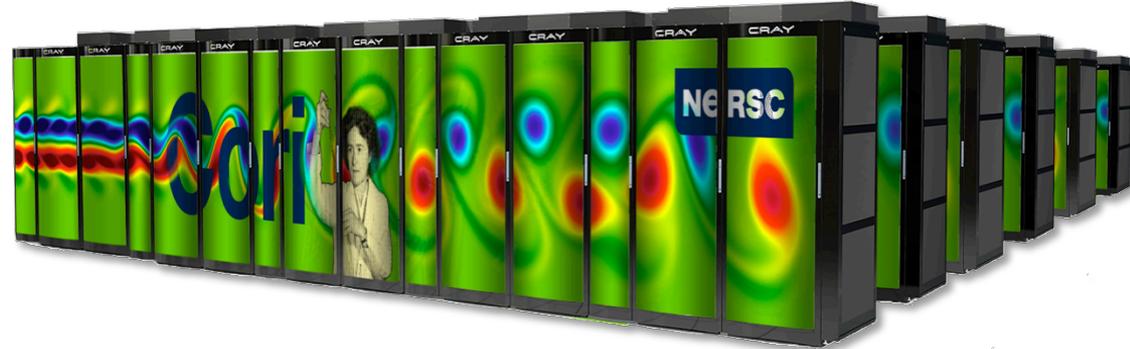
NERSC Senior Science Advisor
High Performance Computing Department Head

National Energy Research Scientific Computing Center



Cori

9,300 Intel Xeon Phi "KNL" manycore nodes
2,000 Intel Xeon "Haswell" nodes
700,000 processor cores, 1.2 PB memory
Cray XC40 / Aries Dragonfly interconnect
30 PB Lustre Cray Sonexion scratch FS
1.5 PB Burst Buffer



#5 on list of Top 500 supercomputers in the world



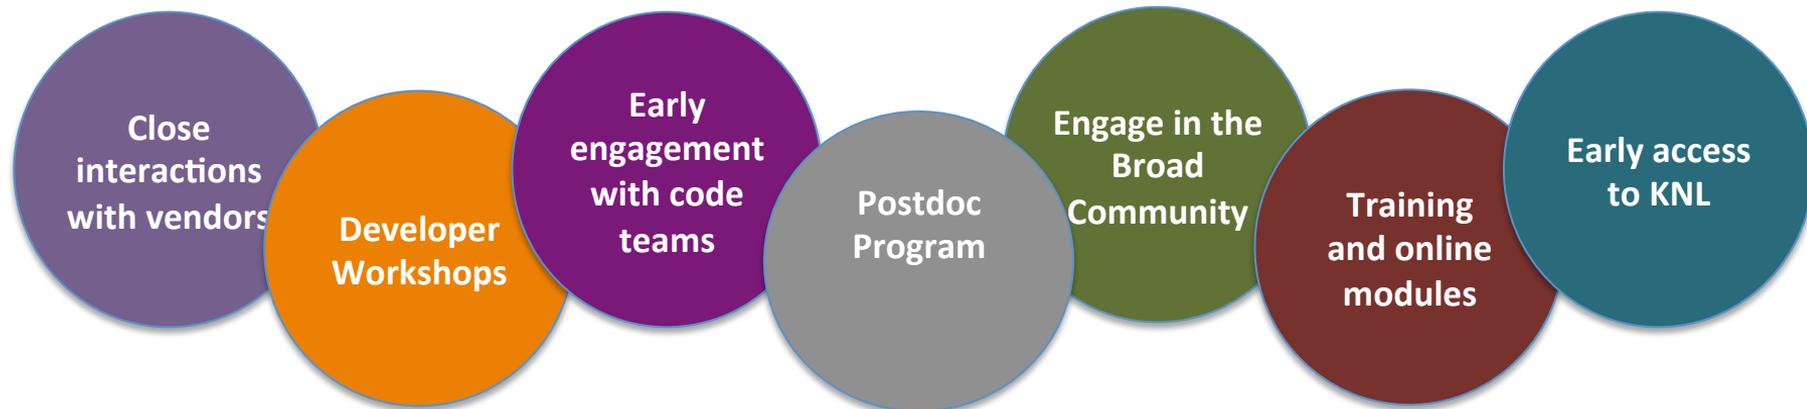
Edison

5,560 Ivy Bridge Nodes / 24 cores/node
133 K cores, 64 GB memory/node
Cray XC30 / Aries Dragonfly interconnect
6 PB Lustre Cray Sonexion scratch FS

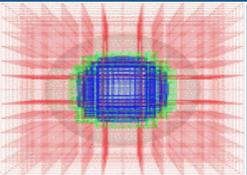
Goal: Prepare DOE Office of Science users for Cori's manycore CPUs

Partner closely with ~20 application teams and apply lessons learned to broad NERSC user community

NESAP for Data is now underway as well



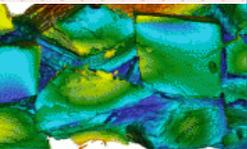
Application Readiness: NESAP



ASCR

Almgren (LBNL)
Trebotich (LBNL)

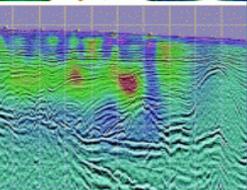
BoxLib
Chombo-crunch



HEP

Vay (LBNL)
Toussaint(Arizona)
Habib (ANL)

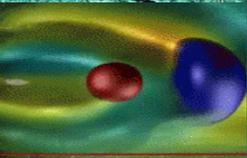
WARP & IMPACT
MILC
HACC



NP

Maris (Iowa St.)
Joo (JLAB)
Christ (Columbia/BNL)

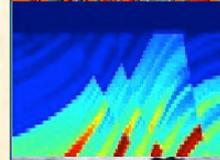
MFDn
Chroma
DWF/HISQ



BES

Kent (ORNL)
Deslippe (LBNL)
Chelikowsky (UT)
Bylaska (PNNL)
Newman (LBNL)

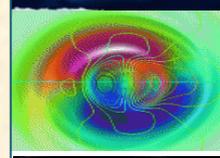
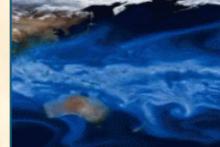
Quantum Espresso
BerkeleyGW
PARSEC
NWChem
EMGeo



BER

Smith (ORNL)
Yelick (LBNL)
Ringler (LANL)
Johansen (LBNL)
Dennis (NCAR)

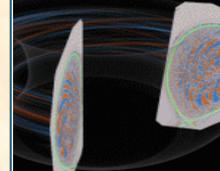
Gromacs
Mercurious
MPAS-O
ACME
CESM



FES

Jardin (PPPL)
Chang (PPPL)

M3D
XGC1



Call for proposals in 2016 and six projects selected

- **Dark Energy Spectroscopic Instrument Codes; Stephen Bailey, Berkeley Lab (HEP)**
- **Union of Intersections Framework; Kris Bouchard, Berkeley Lab (BER)**
- **Cosmic Microwave Background Codes (TOAST); Julian Borrill, Berkeley Lab (HEP)**
- **ATLAS Simulation/Analysis Code; Steve Farrell, Berkeley Lab (HEP)**
- **Tomographic Reconstruction; Doga Gursoy, Argonne (BES)**
- **CMS Offline Reconstruction Code; Dirk Hufnagel, FermiLab (HEP)**

NESAP teams are testing code and running science problems up to full scale on Cori KNL nodes with 24 hours time limit.

All NERSC users are enabled to run on up to 512 nodes for two hours at a time.

Users can gain access to the full system upon request and demonstration of code readiness for KNL.

Usage is not being charged through June 30, 2017.

System still has a few (software) bugs, which users are helping us discover and solve.

NERSC AY 2017 Allocations Forecast



System	"NERSC Hour" Charge per Node Hour	Nodes in System	~Hours in a Year	Overall System Availability Estimate	~Total NERSC Hours for AY2017 (M)	DOE Prod NERSC Hours (M) (80%)	ALCC NERSC Hours (M) (10%)	Directors Reserve NERSC Hours (M) (10%)
Edison	48	5576	8760	.85	2,000	1,600	200	200
Cori P1	80	2000	8760	.85	1,000	800	100	100
Cori P2 (6 months)	96*	9300	8760	.40 (6 months)	3,000†	2,400†	300‡	300†
2017					6,000	4,800	600	600
2016					3,000	2,400	300	300

Each Office/Reserve will have 2X the number of NERSC Hours in 2017 compared to 2016

* - Estimate, may adjust once we measure application performance on system
 † - Supplemental allocation in Spring 2017
 ‡ - Applies to 2017-18 ALCC allocation cycle

Additional hours – equal again to initial allocation – placed in reserves on June 1, 2017 (Discuss)

~2.4 billion additional hours for use on Cori KNL

NERSC hours are fungible among Cori KNL, Cori Haswell, Edison, so codes must run on Cori KNL to use additional hours and therefore they must be ready to run on the architecture

NERSC will send pre-production Cori KNL usage report by repo to DOE managers on June 1, 2017 (Discuss)

Allocation managers transfer resources to selected repos after June 1, 2017 (annotate xfer with "KNL")

The currency for charging at NERSC is the the "NERSC Hour".
1 NERSC Hour ~ 1 core hour on Hopper (retired 2015)

Node Type	Cores per Node	Charge per Node per Hour
Cori KNL	68	96*
Cori Haswell	32	80
Edison	24	48
Hopper (retired)	24	24

* - Production value TBD

Office	KNL NERSC Hours (M)
ASCR	196 (includes app readiness repo)
BER	162
BES	101
FES	7
HEP	96
NP	517

KNL Node Usage AY2017



Project	PI	NE Hrs
Domain Wall Fermions and Highly Improved Staggered Quarks for Lattice QCD (DWF)	Christ (Columbia)	174 M
Quantum Chromodynamics with four flavors of dynamical quarks (MILC)	Toussaint (U Arizona)	69 M
Lattice QCD Monte Carlo Calculation of Hadronic Structure and Spectroscopy (Chroma)	Liu (U Kentucky)	35 M
Synthetic Spectra of Astrophysical Objects	Baron (U Oklahoma)	24 M
Accelerated Climate Modeling for Energy (ACME)	Leung (PNL)	22 M
Molecular Dynamics Simulations of Protein Dynamics and Lignocellulosic Biomass (GROMACS)	Smith (ORNL)	21 M
Center for Computational Study of Excited-State Phenomena in Energy Materials (BerkeleyGW)	Deslippe (LBNL)	20 M
Theory of nanostructured materials (BerkeleyGW, Quantum Espresso)	Neaton (LBNL)	18 M
Large Scale 3D Geophysical Inversion & Imaging (EMGeo)	Newman (LBNL)	10 M

2016 in Review



Program	Target (M)	NERSC Hours Used	Hours Charged
DOE Production	2,400	2,588	2,284
ALCC	300	291	226
DDR	300	69.7	66.9
TOTAL	3,000	2,949	2,576
Startup	-	3.5	1.7
Education	-	1.3	1.2

Office / Program Usage 2016



Office	Initial Allocation	Final Allocation*	Hours Used	Hours Charged	Scavenger Hours*	Balance
ASCR	142	151	135	131	4	20
BER	455	493	431	459	7.4	35
BES	712	758	715	662	67.7	97
FES	429	449	489	401	2.2	49
HEP	366	371	468	353	2.5	17
NP	264	278	372	265	21.6	13
SBIR	20.4	21.1	11	12.6	0	8.5
TOTAL	2,400	2,528	2,588	2,284	100	224

*Refunds

** Includes ALCC and DDR

Allocation Schedule



Date	Item	
January 9, 2017	AY 2017 begins	
April 12, 2017	Allocation reduction	If usage <10%, return 25% of balance to reserve
June 1, 2017 (discuss)	Cori KNL time transferred to reserves	
July 1, 2017	Cori KNL charging begins; New ALCC year begins	
September 18, 2017	ERCAP application deadline	
October 11, 2017	Allocation reduction	<40%, return 25%; <20%, return 75%, <10%, return 90%
December 6, 2017	Allocation reduction	If no usage, close and return balance to reserve
December 8, 2017	Award decision announcement	

Large job discount created extra allocation: ~500 M hours equivalent

Cori downtime in September for KNL node integration

Placed huge pressure on Edison, resulting in long wait times

Payoff will be 2X hours in 2017 and 3X in 2018.

Edison Queue Backlog

Queue Backlog Over Time



Cori KNL nodes are free through July 1, 2017

Charging on Cori KNL nodes starts July 1, 2017

Allocation augmentations placed in reserves on June 1

Program managers will distribute additional time

Codes need to be ready to use the Xeon Phi and program managers need to consider readiness in allocation decisions

NERSC will send KNL usage report and other advice

Cross cut meeting March 9-10 in Tyson's Corner

<http://www.ornl.gov/crosscut2017/>

Goal: integrated summary report of six requirements reviews

- Identify common cross-cutting needs
- Find areas of singular importance to each office
- Discuss a collaborative path forward

Rename startup to "prototype"?

NERSC 9 procurement is underway for 2020 system

Edison will remain through sometime in 2019

DOE Production allocations will not decrease in 2019

Working with ECP in various areas